REDACTED

Group Discovery

   REDACTED

TABLE OF CONTENTS

Groups are collections of peers with a mutual agreement to distribute one or more types of information. Groups are characterized by a *group friendly name* **NAME$_G$** and a *group ID* **CID$_G$**.

Group discovery can be divided into three categories.
   1) Discovery of a private group: You're given CID$_G$.
   2) Discovery of a public group: You choose a name to search for.
   3) Enumeration of a group hierarchy. You want to browse a list of groups.

PNRP can be used for (1) and (2). Group-based mechanisms should be used for (3).

PNRP ID's

PNRP uses peer ID's to uniquely identify peers. Category ID's may be coupled with peer ID's to uniquely identify an instance of a category. Category ID's are not necessarily unique.

Even though CID$_G$ is unique, it should only be used as a PNRP category ID, not a peer ID. Multiple members of the group may register in that category to make discovery easier.

It may be useful to have a non-unique category ID based on a friendly name. NAME$_G$ may be transformed to a category ID **CID$_{NAME}$** by performing a hash of the group friendly name with the seed "GROUP". The seed helps reduce collisions with other types of registered categories with the same name, but different type.

Each peer X is has peer ID **ID$_{P_3}$**.

There are five utility peer ID's. These ID's may be used by actual peers, but are the basis of many important calculations.

1. ID_{MIN} is the lowest possible peer ID.
2. ID_{MAX} is the highest possible peer ID
3. ID_{MID} is $(ID_{MIN} + ID_{MAX}) / 2$.
4. ID_{HMID} is $(ID_{MID} + ID_{MAX}) / 2$.
5. ID_{LMID} is $(ID_{MIN} + ID_{MID}) / 2$.

Uh, the number space is circular. It is hard to assess a "lowest possible peer" in abstract. OTOH, within a group, the fixed prefix makes the number space sequential.

Category ID's are concatenated to peer ID's to generate a categorized peer ID, **CID:ID**. Peer ID's and categorized peer ID's may be registered with PNRP, and searched for through PNRP.

Private group discovery

Private groups are not intended to be easily discoverable. Those invited to join the group must know CID_G to join. If the group has group contacts, each contact Cx registers $CID_G:ID_{Cx}$ with PNRP. If the group does not use contacts, then each group member Px registers $CID_G:ID_{Px}$ with PNRP.

Private group discovery follows this procedure.
1. Potential member P learns CID_G out of band (email, IM, etc.).
2. P searches for $CID_G:ID_{MID}$ using PNRP.
3. If the result is NOT category CID_G
   a. If the return result is for a category ID less than CID_G, search for $CID_G:ID_{HMID}$; else search for $CID_G:ID_{LMID}$.
   b. If the CID of the second search result is still not CID_G, group discovery fails. Go to <u>DONE</u>.
4. Contact the group member described in the search result and join the group.
5. If the group does not use contacts, or if the group uses contacts and you are a contact, register $CID_G:ID_P$ with PNRP.
6. <u>DONE</u>

Ideally a search for $CID_G:ID_{MID}$ will find a member of the desired group. We add a second search to improve robustness for non-ideal PNRP networks.

REDACTED

REDACTED

REDACTED

REDACTED

REDACTED